Survey of Software on Clusters

1. Communication Protocols
   (a) MPI
      i. MPICH
      Features: Free. The MPI-2 standard C++ bindings are available for the MPI-1 functions. Both Fortran 77 and Fortran 90 bindings, including both mpif.h and an MPI module.
      Source: `http://www-unix.mcs.anl.gov/mpi/mpich/`
      Prereqs: rsh/ssh
      Known cluster users: Cyborg from Drexel, ORNL, Hive at GSFC

      ii. LAM/MPI
      Features: Free, Many MPI-2 features are also included in LAM, such as: dynamic process spawning, MPI client/server functionality, one-sided communication, C++ bindings, and MPI I-O. LAM also includes features for fault tolerance.
      Source:`http://www.lam-mpi.org/`
      Prereqs: rsh/ssh
      Known cluster users: ORNL, Hive at GSFC, AENEAS at UC Irvine

   Some reserchers have attempted to compare the performance of MPICH and LAM. They found LAM to be provide better performance and scalability over MPICH [3].

   (b) PVM: From ORNL. Overall objective is to use a heterogeneous collection of computers cooperatively.

   MPI provides extensive messaging support compared to PVM. MPI performance is found to be better than that of PVM [1, 2]. PVM is expected to perform better over a heterogeneous collection of computers.

2. Compilers, Debuggers

(a) C: gcc

(b) Fortran

i. Lahey Fortran
Source: `http://www.lahey.com/`
Includes a command line FDB debugger

ii. Absoft Fortran
Source: `http://www.absoft.com/`
Includes source level debugger with graphical interfaces. Includes shared memory support, widely used on clusters.

iii. Portland Group
Source: `http://www.pgroup.com/`
Includes debugger, profilier

iv. NAGWare f90compiler
Source: `http://www.nag.co.uk/nagware/np.asp`
Not available for 32 bit architectures.

v. XMPI: A Run/Debug GUI for MPI
Source: `http://www.lam-mpi.org/software/xmpi/`

(c) Pricing

Table 1: Pricing Information for various fortran compilers

	No of Users	Upto 16 CPUs	Upto 64 CPUs	Upto 256 CPUs
Lahey Fortran PRO	2	($2085, $1490)	($3130, $2250)	($4690,$3370)
(Commercial,	5	($3990,$2865)	($5985,$4300)	($8980,$6450)
Educational)	10	($7500,$5385)	($11260,$8080)	($16890, $12120)
	2	($1899, $1465)	$3499, $2905)	($6299,$4725)
Absoft	5	($2695,$2019)	($4895,$4069)	($8895, $6619)
	10	($4319, $3230)	($7819, $6510)	($14219, $10590)
portland group	2	($2499, $1999)	($5249,$4199)	($10498, $8398)
	5	($3599,$2879)	($7599,$7599)	($15198,$15918)
	10	($5849, $4679)	($12349, $12470)	-

3. Profiling [4].

(a) NupShot/upshot/Jumpshot
Source: `http://www-fp.mcs.anl.gov/~lusk/upshot/`
Language: Language - independent
Prereqs: Tcl/Tk or Java
Tested platforms: SGI, IBMSP, Solaris
Works for MPI

(b) Paradyn: Free
Source : `http://www.cs.wisc.edu/paradyn/`
Languages: Fortran 90, C
Platforms: Solaris, Linux, Windows, AIX
Works for MPI and PVM

(c) VAMPIR:
Vampirtrace is an instrumented MPI library.
Source: `http://www.pallas.de/pages/vampir.htm` Languages: Language Independent
Platforms: Linux, Solaris, IBM SP, etc.
Pricing : Vampir - Single system : $1700
Vampirtrace 1-8 procs : $2600
32 procs : $5200
128 procs : $10,400
Works for MPI and PVM

4. Queuing Systems

(a) DQS - distributed queueing system:
Free
Source: `http://www.scri.fsu.edu/~pasko/dqs.html`
The Distributed Queueing System is designed as a management tool to aid in computational resource distribution across a network. DQS provides architecture transparency for both users and administration across a heterogeneous environment, allowing for seamless interaction for multiple architectures.

(b) Condor
Free
Source: `http://www.cs.wisc.edu/condor/`
Condor is a specialized workload management system for compute-intensive

jobs. Like other full-featured batch systems, Condor provides a job queueing mechanism, scheduling policy, priority scheme, resource monitoring, and resource management. Users submit their serial or parallel jobs to Condor, Condor places them into a queue, chooses when and where to run the jobs based upon a policy, carefully monitors their progress, and ultimately informs the user upon completion.

References

[1] J. Kitowski, K. Boryczko, and J. Moscinski. Comparison of PVM and MPI performance in short-range molecular dynamics simulation. In M. Bubak, J. Dongarra, and J. Wasneiwski, editors, *Proceedings of the 4th European* PVM/MPI *User's Group Meeting*, Lecture Notes in Computer Science 1332, pages 11–16, Cracow, Poland, 1997.

[2] S. Markus, S. B. Kim, K. Pantazapoulos, A. L. Ocken, E. N. Houstis, P. Wu, S. Weerawarana, and D. Maharry. Performance comparison of MPI implementations using parallel ELLPACK PSE. In *Proceedings of the Second* MPI *Developer's Conference*, pages 162–169. IEEE Computer Society Press, 1996.

[3] http://www-hpcc.astro.washington.edu/faculty/trq/brandon/perform.html. Comparison of MPICH and LAM.

[4] http://www.cs.utk.edu/ browne/perftools-review/. Review of performance analysis toolsfor mpi parallel programs.